
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Basavaiah et al.

Attorney Docket No. ANDIP037

Application No.: 10/726,269

Examiner: Unelus, Ernest

Filed: December 1, 2003

Group: 2181

Title: APPARATUS AND METHOD FOR
PERFORMING FAST FIBRE CHANNEL
WRITE OPERATIONS OVER RELATIVELY
HIGH LATENCY NETWORKS

Confirmation No. 3368

CERTIFICATE OF EFS-WEB TRANSMISSION

I hereby certify that this correspondence is being transmitted electronically through EFS-WEB to the Commissioner for Patents, P.O. Box 1450 Alexandria, VA 22313-1450 on March 7, 2007.

Signed: /Linda L. Quintana/
Linda L. Quintana

PRE-APPEAL BRIEF REQUEST FOR REVIEW

Mail Stop AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

The Applicants' hereby request review of the rejections in the above-identified application. This request is being filed with a Notice of Appeal. The review is requested for the reasons stated below.

STATUS OF CLAIMS

This appeal is from the Final Office Action dated September 29, 2006. In the Final Office Action, the Examiner rejected claims 1-23 under 35 USC 102(e) as anticipated by Mullendore (US Patent Publication 2003/0185154). The Applicants' filed a response after final on November 10, 2006. The Examiner maintained the rejection in an Advisory Action dated December 6, 2006. Claims 1-23 remain at issue and are the subject of this pre-appeal brief.

BACKGROUND TO THE PRESENT INVENTION

Within a Storage Array Network (SAN), a Host may write to a target storage device using a write command. The Host will issue a write (Wr) command defining a certain amount of data to be written. The command travels across the network, from Switch to Switch, until it reaches the target. In reply, the target responds with a Xfer ready command which defines the amount of data which the target may accept. When the Host receives the Xfer ready command, it transfers the data to be written in units up to the maximum transfer unit (MTU) of the network. If, for some reason, the Host does not receive the status command after a predetermined period of time, it is assumed that a problem occurred. The Host may subsequently issue another write command.

With two SANs coupled together by a high latency inter-SAN network, a significant amount of time may lapse between the time the initial Wr command is issued by a Host in the first SAN and the Xfer ready is received by the Host from the target device in the second SAN due. During this time, the Host is idle and must wait before issuing the data transfer commands to transfer the data to the Host.

THE PRESENT INVENTION

The present invention relates to a fast write operation between a Host (H1) in one SAN and a target device (T1) in a second SAN where the two SANs are inter-connected by a high latency network. As best illustrated in Figures 1 and 4 of the present application, the sequence of the fast write of the present invention includes the following:

- a. Host H1 initiates the fast write operation by issuing a SCSI write command (Wr: OX_ID = 1 RX_ID = 00xffff, Size = 10MB). The command defines the originating exchange identifier as 1 (OX_ID = 1). The receiving exchange identifier RX_ID, however, is “uninitialized” and is set to a default value of “0xffff”. The write command also specifies the amount of data to be written, which in this example, is 10 megabytes (MB).
- b. Upon receipt, the switch SW1 initializes the receiving exchange identifier RX_ID. In this example, the RX_ID is initialized to 10. Assuming the switch SW1 has sufficient storage space to buffer the data, the switch SW1 sends a Transfer Ready command (Xrdy: OX_ID = 1, RX_ID = 10, Size = 10 MG) to the Host H1. All subsequent commands or frames between the Host and switch SW1, and vice

versa, associated with this SCSI write operation are define the OX_ID = 1 and the RX_ID =10.

- c. The initiating switch SW1 uses the OX_ID to keep track of the transaction. Consequently, the switch SW1 changes the OX_ID provided by the initiating Host H1. In this example, the switch SW1 changes the OX_ID value to 10. The switch SW1 then forwards the write command to the target T1 with the RX_ID value remaining uninitialized (Wr: OX_ID = 10, RX_ID = Oxffff, Size = 10MB). All communication between the first switch SW1 and the target involving this write operation thereafter includes an OX_ID = 10 and RX_ID = Oxffff. The initiating switch SW1 uses the OX_ID value as a handle or pointer into a session table 36 maintained at switch SW1. The table includes an entry that includes information regarding the session that is accessed by the RX_ID handle.
- d. When the second switch SW2 receives the write command, it initializes an exchange identifier entry in the sessions table 38 and it immediately forwards the command to the target T1 provided the switch SW2 has sufficient buffer space. If it does not have sufficient space, then a SCSI busy status is sent back to the initiating host H1.
- e. If the target T1 is ready to receive the data, it sends a Transfer Ready command back to the switch SW2. The target designates an RX_ID value for the write transaction. In this case, the target designates an RX_ID value of 50. The Transfer Ready command received by the switch SW2 therefore appears as (Xrdy: OX_ID = 10, RX_ID = 50, size = 10MB). All subsequent communications between the switch SW2 and the target T1 involving this transaction include OX_ID value of 10 and an RX_ID value of 50. The switch SW2 also maintains a sessions ID table 38. Upon receipt of the Transfer Ready command, the switch SW2 inserts a RX_ID = 50 value into the table. The switch SW2 uses the modified OX_ID = 10 value as a handle or pointer into a sessions ID table 38. The target switch SW2 uses the OX_ID value as a handle or pointer for this session between in session table 38. The table includes an entry that includes the information regarding the session such as the target RX_ID.
- f. If the second switch SW2 receives the data frames (Wdata: OX_ID = 10, RX_ID = Oxffff) from the first switch SW1 before the Transfer Ready command from the

target T1, then the second switch SW2 buffers the data. When the Transfer Ready command is received, the data frame(s) are then forwarded to the target T1. On the other hand, if the data frames arrive after the Transfer Ready command, the data frames are immediately forwarded to the target T1.

- g. When all the data has been transferred, the target T1 generates a Status command (Status: OX_ID=10, RX_ID=50). The second switch SW2 modifies the RX_ID = to equal 0xffff and forwards the status command to the switch SW1. The switch SW1 in turn changes the RX_ID = 10 and sends the status command to the Host H1 to complete the fast write operation.

THE CLAIMS OF THE PRESENT INVENTION

The claims of the present invention relate to an apparatus for implementing a fast write operation in a Switch used in a SAN. The Switch includes a port configured to receive a write command defining an initiating Host and a target, and a trapping mechanism to trap the write command if it designates a predetermined Host_ID and a predetermined target_ID. The Switch also includes a processor configured to process the trapped write commands by **modifying either the OX_ID or the RX_ID** of the command.

THE REJECTION

While the Applicants' acknowledge that Mullendore describes a fast write operation, the reference fails to teach or suggest a fast write where either the OX_ID or the RX_ID of the write command is modified. On the contrary, the fast write of Mullendore involves the Switches issuing a Ready To Transfer or "RTT" message (i.e., which is the same as the Xfer ready command used by the Applicants') to the initiating Host on behalf of the target, in order to enhance performance by causing the initiator to "fill the pipe" with write data. See paragraph [0060].

With reference to Figure 4, Mullendore describes the prior art fast write in detail in Paragraph [0061]. When the Fast Write is enabled, the Switch 150 does not issue a Ready-to-Transfer (RTT) message unless the Switch has buffer resources available to buffer the data of the entire write operation or some threshold amount (e.g., 256KB). If the original write command transfers less than the threshold amount, or the buffer resources are available for all the data, then the RTT message is immediately issued. If the original write command transfers more than the

available buffer space, then the RTT message is generated only when the buffer space becomes available. If the write command transfers more data than the threshold, then multiple threshold sized RTT messages are generated as buffer space becomes available.

In paragraphs [0069] through [0072] and Figures 6 and 7, Mullendore describes the differences between a standard write operation and a fast write operation.

A standard write operation (i.e., 1MB) is initiated by initiator 235. In response, the target 245 responds with an RTT message, which defines the amount of available buffer space (i.e., 512 KB). The initiator 235 responds to the RTT by transferring 512KB to the target. This process is repeated until the 1 MB of data is transferred from the initiator 235 to the target 245. See Figure 6 and paragraph [0070].

With the Fast write operation, the switch 250 near the initiator 235 issues the RTT requesting the entire 1MB of data in response to the write operation. The 1 MB of data is then forwarded and stored in the far end switch 240, near the target 245. A series of RTT messages are then transmitted between the target 245 and the far end switch 240 until the 1 MB is written to the target.

Absolutely nowhere in the Mullendore reference are the OX_ID and/or the RX_ID of a write command mentioned, let alone altered by a Switch.

Applicants' believe that all pending claims have been amended and the case is now in a condition for allowance. The Applicants' respectfully request a Notice of Allowance for this application from the Examiner. Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

Respectfully submitted,
BEYER WEAVER LLP

/James W. Rose/
James W. Rose
Reg. No. 34,239

P.O. Box 70250
Oakland, CA 94612-0250
(408) 255-8001